U.S. HOUSE OF REPRESENTATIVES
COMMITTEE ON SCIENCE

HEARING CHARTER

Supercomputing: Is the U.S. on the Right Path?

Wednesday, July 16, 2003
10:00 a.m. - 12:00 Noon
2318 Rayburn House Office Building

1. Purpose

On Wednesday, July 16, 2003, the House Science Committee will hold a hearing to examine whether the United States is losing ground to foreign competitors in the production and use of supercomputers[1] and whether federal agencies' proposed paths for advancing our supercomputing capabilities are adequate to maintain or regain the U.S. lead.

2. Witnesses

Dr. Raymond L. Orbach is the Director of the Office of Science at the Department of Energy. Prior to joining the Department, Dr. Orbach was Chancellor of the University of California at Riverside.

Dr. Peter A. Freeman is Assistant Director for the Computer and Information Science and Engineering Directorate (CISE) at the National Science Foundation (NSF). Prior to joining NSF in 2002, he was professor and founding Dean of the College of Computing at Georgia Institute of Technology.

Dr. Daniel A. Reed is the Director of the National Center for Supercomputing Applications (NCSA) at the University of Illinois at Urbana-Champaign. NCSA is the leader of one of NSF's two university-based centers for high-performance computing. Dr. Reed is also the Director of the National Computational Science Alliance and is a principal investigator in the National Science Foundation's TeraGrid project. Earlier this year, Dr. Reed was appointed to the President's Information Technology Advisory Committee (PITAC).

Mr. Vincent Scarafino is the Manager of Numerically Intensive Computing at Ford Motor Company, where he focuses on providing flexible and reliable supercomputer resources for Ford's vehicle product development, including vehicle design and safety analysis.

3. Overarching Questions

The hearing will address the following overarching questions:

1. Is the U.S. losing its leadership position in supercomputing? Do the available supercomputers allow United States science and industry to be competitive internationally? Are federal efforts appropriately

---

[1] Supercomputing is also referred to as high-performance computing, high-end computing, and sometimes advanced scientific computing.

targeted to deal with this challenge?

2. Are federal agencies pursuing conflicting supercomputing programs? What can be done to ensure that federal agencies pursue a coordinated policy for providing supercomputing to meet the future needs for science, industry, and national defense?

3. Is the National Science Foundation moving away from the policies and programs that in the past have provided broad national access to advanced supercomputers?

4. Can the U.S. fulfill its scientific and defense supercomputing needs if it continues to rely on machines designed for mass-market commercial applications?

4. Brief Overview

- High-performance computers (also called supercomputers) are an essential component of U.S. scientific, industrial, and military competitiveness. However, the fastest and most efficient supercomputer in the world today is in Japan, not the U.S. Some experts claim that Japan was able to produce a computer so far ahead of the American machines because the U.S. had taken an overly cautious or conventional approach for developing new high-performance computing capabilities.

- Users of high-performance computing are spread throughout government, industry, and academia, and different high-performance computing applications are better suited to different types of machines. As the U.S. works to develop new high-performance computing capabilities, extraordinary coordination among agencies and between government and industry will be required to ensure that creative new capabilities are developed efficiently and that all of the scientific, governmental, and industrial users have access to the high-performance computing hardware and software best suited to their applications.

- The National Science Foundation (NSF) currently provides support for three supercomputing centers: the San Diego Supercomputer Center, the National Center for Supercomputing Applications at Urbana-Champaign in Illinois, and the Pittsburgh Supercomputing Center. These centers, along with their partners at other universities, are the primary source of high-performance computing for researchers in many fields of science. Currently, support for these centers beyond fiscal year 2004 is uncertain, and in the past few years NSF has been increasing its investment in a nationwide computing grid, in which fast connections are built between many computers to allow for certain types of high-performance scientific computing and advanced communications and data management. It is not clear whether this “grid computing” approach will provide the high-performance computing capabilities needed in all the scientific fields that currently rely on the NSF supercomputing centers.

- At the Department of Energy, there are two programs aimed at advancing high-performance computing capabilities. One, in the National Nuclear Security Administration (NNSA), is the continuation of a long-term effort to provide supercomputers to be used for modeling nuclear weapons effects; these simulations are particularly important in light of existing bans on nuclear weapon testing. In the other program, the Office of Science is now proposing to supplement its current advanced scientific computing activities with a new effort designed to create the world’s fastest supercomputers.

5. Current Issues

Is the U.S. Competitive?

Japan's Earth Simulator is designed to perform simulations of the global environment that allow researchers to study scientific questions related to climate, weather, and earthquakes. It was built by NEC for the Japanese government at a cost of at least $350 million and has been the fastest computer in the world since it began running in March 2002. When the first measures of its speed were performed in April 2002, researchers determined that the Earth Simulator was almost five times faster than the former record holder, the ASCI White System at Lawrence Livermore National Laboratory, and also used the machine's computing power significantly more efficiently.[2]

This new development caused a great deal of soul-searching in the high-performance computing community about the U.S. approach to developing new capabilities and the emphasis on using commercially available (not specialized or custom-made) components. Going forward, it is not clear whether or not such a commodity-based approach will allow the U.S. high-performance computing industry to remain competitive. It is also unclear if the new machines produced by this approach will be able provide American academic, industrial, and governmental users with the high-performance computing capabilities they need to remain the best in the world in all critical applications.

Will All Users Be Served?

Users of high-performance computing are spread throughout government, industry, and academia. Different high-performance computing applications are better suited to different types of machines. For example, weather modeling and simulations of nuclear weapons require many closely-related calculations, so machines for these applications must have components that communicate with each other quickly and often. Other applications, such as simulations of how proteins fold, can be efficiently performed with a more distributed approach on machines in which each component tackles a small piece of the problem and works in relative isolation. In the U.S., the major producers of high-performance computers include IBM, Hewlett-Packard, and Silicon Graphics, Inc., whose products lean toward the more distributed approach, and Cray, whose products are more suited to problems that require the performance of closely-related calculations. The Japanese (NEC, Fujitsu, and Hitachi), also produce this sort of machine. The concern is that the U.S. on the whole has moved away from developing and manufacturing the machines needed for problems with closely-related calculations because the more distributed machines have a bigger commercial market. The Japanese have been filling this gap, but the gap could still impact the access of American scientists to the types of supercomputers that they need for certain important research problems.

Responsibility for providing high-performance computing capabilities to existing users and for developing new capabilities is distributed among 11 different federal agencies and offices and relies heavily on industry for development and production. In this environment, extraordinary amounts of coordination are needed to ensure that new capabilities are developed efficiently and that the most appropriate kinds of hardware and software are available to the relevant users—coordination among agencies and between government and industry, as well as cooperation among universities and hardware and software companies. The results of an ongoing interagency effort to produce a coherent high-performance computing roadmap and the influence this roadmap has on agencies' programs will be the first test.

Where are the DOE Office of Science and the NSF Programs Headed?

Both NSF and the DOE Office of Science are moving ahead in significant new directions. At NSF, no plans have been announced to continue the Partnerships for Advanced Computational

[2] For the U.S. supercomputers, typical scientific applications usually only are able to utilize 5-10 percent of the theoretical maximum computing power, while the design of the Earth Simulator makes 30-50 percent of its power accessible to the majority of typical scientific applications.

Infrastructure program, which supports the supercomputer centers, beyond fiscal year 2004. In addition, a proposed reorganization of NSF's Computer and Information Sciences and Engineering Directorate was announced on July 9 that includes a merging of the Advanced Computational Infrastructure program (which includes the support for the supercomputer centers) and the Advanced Networking Infrastructure program (which supports efforts on grid computing—an alternative approach to high-performance computing). Some scientists have expressed concerns that NSF may be reducing its commitment to providing researchers with a broad range of supercomputing capabilities and instead focusing its attention on grid computing and other distributed approaches.

For the DOE Office of Science, the fiscal year 2004 budget request proposes a new effort in next-generation computer architecture to identify and address major bottlenecks in the performance of existing and planned DOE science applications. In addition, the July 8 mark-up of the House Energy and Water Development Appropriations Subcommittee sets funding for the Advanced Scientific Computing Research initiative at \$213.5 million, an increase of \$40 million over the request and \$46 million over the previous year. Decisions about the future directions for high-performance computing at NSF and DOE Office of Science are clearly being made now.

The White House has an interagency effort underway, the High End Computing Revitalization Task Force (HECRTF), which is supposed to result in the agencies' submitting coordinated budget requests in this area for fiscal year 2005.

6. Background

What is High-Performance Computing? High-performance computing—also called supercomputing, high-end computing, and sometimes advanced scientific computing—is a phrase used to describe machines or groups of machines that can perform very complex computations very quickly. These machines are used to solve complicated and challenging scientific and engineering problems or manage large amounts of data. There is no set definition of how fast a computer must be to be "high-performance" or "super," as the relevant technologies improve so quickly that the high-performance computing achievements of a few years ago could be handled now by today's desktops. Currently, the fastest supercomputers are able to perform trillions of calculations per second.

What is High-Performance Computing Used For? High-performance computing is needed for a variety of scientific, industrial, and national defense applications. Most often, these machines are used to simulate a physical system that is difficult to study experimentally. The goal can be to use the simulation as an alternative to actual experiments (e.g. for nuclear weapon testing and climate modeling), as a way to test our understanding of a system (e.g. for particle physics and astrophysics), or as a way to increase the efficiency of future experiments or product design processes (e.g. for development of new industrial materials or fusion reactors). Other major uses for supercomputers include performing massive complex mathematical calculations (e.g. for cryptanalysis) or managing massive amounts of data (e.g. for government personnel databases).

> *Scientific Applications:* There are a rich variety of scientific problems being tackled using high-performance computing. Large-scale climate modeling is used to examine possible causes and future scenarios related to global warming. In biology and biomedical sciences, researchers perform simulations of protein structure, folding, and interaction dynamics and also model blood flows. Astrophysics model planet formation and supernova, and cosmologists analyze data on light from the early universe. Particle physicists use the ultra-fast computers to perform the complex calculations needed to study quantum chromodynamics and improve our understanding of electrons and quarks, the basic building blocks of all matter. Geologists model the stresses within the earth to study plate tectonics, while civil engineers simulate the impact of earthquakes.

National Defense Applications: There are a number of ways in which high-performance computing is used for national defense applications. The National Security Agency (NSA) is a major user and developer of high-performance computers for executing specialized tasks relevant to cryptanalysis (such as factoring large numbers). The Department of Energy's National Nuclear Security Administration is also a major user and developer of machines to be used for designing and modeling nuclear weapons. Other applications within the Department of Defense include armor penetration modeling, weather forecasting, and aerodynamics modeling. Many of the scientific applications also have direct or future defense applications. For example, computational fluid dynamics studies are also of interest to the military, e.g. for modeling turbulence around aircraft. The importance of high-performance computing in many military areas, including nuclear and conventional weapons design, means that machines that alone or when wired together are capable of superior performance at military tasks are subject to U.S. export controls.

Industrial Applications: Companies use high-performance computing in a variety of ways. The automotive industry uses fast machines to maximize the effectiveness of computer-aided design and engineering. Pixar uses massive computer animation programs to produce films. Pharmaceutical companies simulate chemical interactions to help with drug design. The commercial satellite industry needs to manage huge amounts of data for mapping. Financial companies and other industries use large computers to process immense and unpredictable Web transaction volumes, mine databases for sales patterns or fraud, and measure the risk of complex investment portfolios.

What Types of High-Performance Computers Are There? All of the above examples of high-performance computing applications require very fast machines, but they do not all require the same type of very fast machine. There are a number of different ways to build high-performance computers, and different configurations are better suited to different problems. There are many possible configurations, but they can be roughly divided into two classes: big, single-location machines and distributed collections of many computers (this approach is often called grid computing). Each approach has its benefits – the big machines can be designed for a specific problem and are often faster, while grid computing is attractive in part because by using a multitude of commercially-available computers, the purchase and storage cost is often lower than for a large specialized supercomputer.

Since the late 1990's, the U.S. approach to developing new capabilities has emphasized using commercially available (not specialized) components as much as possible. This emphasis has resulted in an increased focus on grid computing, and, in large machines, has led to a hybrid approach in which companies use commercial processors (whose speed is increasing rapidly anyway) to build the machines and then further speed them up by increasing the number of processors and improving the speed at which information is passed between processors. There are a number of distinctions that can be made among large machines bases on how the processors are connected. The differences relate to how fast and how often the various components of the computer communicate with each other and how calculations are distributed among the components.

Users thus have a number of options for high-performance computing. Each user must take into account all of the pros and cons of the different configurations when he is deciding what sort of machine to use and how to design software to allow that machine to most efficiently solve his problem. For example, some problems, like weather and climate modeling and cryptanalysis, require lots of communication among computer components and large quantities of stored data, while other applications, like large-scale data analysis for high energy physics experiments or bioinformatics projects, can be more efficiently performed on distributed machines each tackling its own piece of the problem in relative isolation.

How Do Government and Industry Provide Existing and New High-Performance Computing Capabilities? The development and production of high-performance computing capabilities requires significant effort by both government and industry. For any of the applications of high-performance computing described above, the users need good hardware (the high-performance machine or group of machines) and good software (programs that allow them to perform their calculations as accurately and efficiently as possible).

The role of government therefore includes (1) funding research on new approaches to building high-performance computing hardware, (2) in some cases, funding the development stage of that hardware (usually through security agencies), (3) purchasing the hardware to be used by researchers at universities and personnel at government agencies, (4) funding research on software and programs to use existing and new high-performance computing capabilities, and (5) supporting research that actually uses the hardware and software. The role of industry is complementary – i.e. it receives funding to do research and development on new hardware and software, and it is the seller of this hardware and software to government agencies, universities, and companies. The primary industries involved in producing high-performance computing capabilities are computer makers (such as IBM, Hewlett-Packard, Silicon Graphics, Inc., and Cray), chip makers (such as Intel), and software designers. Congress has long had concerns about the health of the U.S. supercomputing industry. In 1996, when the National Center for Atmospheric Research, a privately-run, federally-funded research center, tried to order a supercomputer from NEC for climate modeling, Congress blocked the purchase.

Federal High-Performance Computing Programs: In 1991, Congress passed the High Performance Computing Act, establishing an interagency initiative (now called National Information Technology Research and Development (NITRD) programs) and a National Coordination Office for this effort. Currently 11 agencies or offices participate in the high-end computing elements of the NITRD program (See Table 1 in the appendix). The total requested by all 11 agencies in fiscal year 2003 for high end computing was $846.5 million. The largest research and development programs are at the National Science Foundation (NSF), which requested $283.5 million, and the Department of Energy Office of Science, which requested $137.8 million. Other major agency activities (all between $80 and $100 million) are at the National Institutes of Health (NIH), the Defense Advanced Research Projects Agency (DARPA), the National Aeronautics and Space Administration (NASA), and the Department of Energy's National Nuclear Security Administration (NNSA). Different agencies concentrate on serving different user communities and on different stages of hardware and software development and application. (In addition to the research and development-type activities that are counted for the data included in Table 1 and referenced above, many agencies, such as NNSA and the National Oceanic and Atmospheric Administration (NOAA), devote significant funding to the purchase and operation of high-performance computers that perform these agencies' mission-critical applications.[3])

> *National Science Foundation:* The NSF serves a very wide variety of scientific fields within the academic research community, mainly through a series of supercomputing centers, originally established in 1985 and currently funded under the Partnerships for Advanced Computational Infrastructure (PACI) program. The supercomputer centers provide researchers not only with access to high-performance computing capabilities but also with tools and expertise on how best to utilize these resources. The NSF also is supporting the development of the Extensible Terascale Facility (ETF), a nationwide grid of machines that can be used for high-performance computing and advanced communications and data management. Recently, some researchers within the high-performance computing community have expressed concern that NSF may be reducing its commitment to the supercomputer centers and increasing its focus on grid computing and distributed approaches to high-performance computing, such as would be used in the ETF.

[3] For example, in FY 2003 NOAA spent $36 million on supercomputers—$10 million for machines for climate modeling and $26 million for machines for the National Weather Service.

Department of Energy: The Department of Energy has been a major force in advancing high-performance computing for many years, and the unveiling of the fastest computer in the world in Japan in 2002 resulted in serious self-evaluation at the department, followed by a rededication to efforts to enhance U.S. supercomputing capabilities. The Department of Energy has two separate programs focused on both developing and applying high-performance computing. The Advanced Scientific Computing Research (ASCR) program in the Office of Science funds research in applied mathematics (to develop methods to model complex physical and biological systems), in network and computer sciences, and in advanced computing software tools. For fiscal year 2004, the department has proposed a new program on next-generation architectures for high-performance computing. The Accelerated Strategic Computing Initiative (ASCI) is part of the NNSA's efforts to provide advanced simulation and computing technologies for weapons modeling.

DARPA: DARPA traditionally focuses on the development of new hardware, including research into new architectures and early development of new systems. On July 8, DARPA announced that Cray, IBM, and Sun Microsystems had been selected as the three contractor teams for the second phase of the High Productivity Computing Systems program, in which the goal is to provide a new generation of economically viable, scalable, high productivity computing systems for the national security and industrial user communities in the 2009 to 2010 timeframe.

Other Agencies: NIH, NASA, and NOAA are all primarily users of high performance computing. NIH manages and analyzes biomedical data and models biological processes. NOAA uses simulations to do weather forecasting and climate change modeling. NASA has a variety of applications, including atmospheric modeling, aerodynamic simulations, and data analysis and visualization. The National Security Agency (NSA) both develops and uses high-performance computing for a number of applications, including cryptanalysis. As a user, NSA has a significant impact on the high-performance computing market, but due to the classified nature of its work, the size of its contributions to High End Computing Infrastructure and Applications and the amount of funding it uses for actual operation of computers is not included in any of the data.

Interagency Coordination: The National Coordination Office (NCO) coordinates planning, budget, and assessment activities for the Federal Networking and NITRD Program through a number of interagency working groups. The NCO reports to the White House Office of Science and Technology Policy and the National Science and Technology Council. In 2003, NCO is also managing the High End Computing Revitalization Task Force (HECRTF), an interagency effort on the future of U.S. high-performance computing. The HECRTF is tasked with development of a roadmap for the interagency research and development for high-end computing core technologies, a federal high-end computing capacity and accessibility improvement plan, and a discussion of issues relating to federal procurement of high-end computing systems. The product of the HECRTF process is expected to guide future investments in this area, starting with agency budget submissions for fiscal year 2005.

The Role of Industry: Industry plays a critical role in developing and providing high-performance computing capabilities to scientific, industrial, and defense users. Many supercomputers are purchased directly from computer companies like IBM, Hewlett-Packard, Silicon Graphics, Inc., and Cray, and the groups that do build their own high-performance clusters do so from commercially available computers and workstations. Industry is a recipient of federal funding for initial research into new architectures for hardware, for development of new machines, and for production of standard and customized systems for government and universities, but industry also devotes its own funding to support research and development. The research programs do not just benefit the high-performance computing community, as new architectures and faster chips lay the groundwork for better performing computers and processors in all commercial information technology products.

The State of the Art in High-Performance Computing: Twice a year, a list of the 500 fastest supercomputers is compiled; the latest list was released on June 23, 2003 (see Table 2 in the appendix).[4] The Earth Simulator supercomputer, built by NEC and installed last year at the Earth Simulator Center in Yokohama, Japan, continues to hold the top spot as the best performer. It is approximately twice as fast as the second place machine, the ASCI Q system at Los Alamos National Laboratory, built by Hewlett-Packard. Of the top twenty machines, eight are located at various Department of Energy national laboratories and two at U.S. universities,[5] and nine were made by IBM and five by Hewlett-Packard.

7. Witness Questions

The witnesses were asked to address the following questions in their testimony:

Questions for Dr. Raymond L. Orbach

- The Office of Science appears to have embarked on a new effort in next-generation advanced scientific computer architecture that differs from the development path currently pursued by the National Nuclear Security Agency (NNSA), the lead developer for advanced computational capability at the Department of Energy (DOE). Why is the Office of Science taking this approach?
- How is the Office of Science cooperating with the Defense Advanced Research Projects Agency, which supports the development of advanced computers for use by the National Security Agency and other agencies within the Department of Defense?
- To what extent will the Office of Science be guided by the recommendations of the High-End Computing Revitalization Task Force? How will the Office of Science contribute to the Office of Science and Technology Policy plan to revitalize high-end computing?
- To what extent are the advanced computational needs of the scientific community and of the private sector diverging? What is the impact of any divergence on the advanced computing development programs at the Office of Science?

Questions for Dr. Peter A. Freeman

- Some researchers within the computer science community have suggested that the NSF may be reducing its commitment to the supercomputer centers. Is this the case? To what extent does the focus on grid computing represent a move away from providing researchers with access to the most advanced computing equipment?
- What are the National Science Foundation's (NSF's) plans for funding the supercomputer centers beyond fiscal year 2004? To what extent will you be guided by the recommendation of the NSF

[4] The top 500 list is compiled by researchers at the University of Mannheim (Germany), Lawrence Berkeley National Laboratory, and the University of Tennessee and is available on line at http://www.top500.org/. For a machine to be included on this public list, its owners must send information about its configuration and performance to the list-keepers. Therefore, the list is not an entirely comprehensive picture of the high-performance computing world, as classified machines, such as those used by NSA, are not included.

[5] The two university machines are located at the Pittsburgh Supercomputing Center (supported primarily by NSF) and Louisiana State University's Center for Applied Information Technology and Learning. The remaining 12 machines include four in Europe, two in Japan, and one each at the National Oceanic & Atmospheric Administration, the National Center for Atmospheric Research, the Naval Oceanographic Office, and NASA.

Advisory Panel on Cyberinfrastructure to maintain the Partnerships for Advanced Computational Infrastructure, which currently support the supercomputer centers?

- To what extent will NSF be guided by the recommendations of the High-End Computing Revitalization Task Force? How will NSF contribute to the Office of Science and Technology Policy plan to revitalize high-end computing?
- To what extent are the advanced computational needs of the scientific community and of the private sector diverging? What is the impact of any such divergence on the advanced computing programs at NSF?

Questions for Dr. Daniel A. Reed

- Some researchers within the computer science community have suggested that the National Science Foundation (NSF) may be reducing its commitment to provide advanced scientific computational capability to U.S. scientists and engineers. Have you detected any change in policy on the part of NSF?
- What advanced computing capabilities must the federal government provide the academic research community for the government's programs to be considered successful? Are the programs for developing the next-generation of advanced scientific computing that are currently underway at government agencies on track to provide these capabilities? If not, why not?
- For academic scientists and engineers, what is the difference between the advanced scientific computing capabilities provided by NSF and those provided by the Department of Energy?

Questions for Mr. Vincent F. Scarafino

- How does Ford use high-performance computing? How do computing capabilities affect Ford's competitiveness nationally and internationally?
- What does Ford see as the role of the Federal government in advancing high-performance computing capabilities and in making these capabilities accessible to users? Are current agency programs for developing the next-generation of advanced scientific computing adequate to provide these capabilities? If not, why not?
- Is the U.S. government cooperating appropriately with the private sector on high-performance computing, and is the level of cooperation adequate to sustain leadership and meet scientific and industrial needs?
- To what extent are the advanced computational needs of the scientific community and of the private sector diverging? What is the impact of any divergence on Ford's access to advanced computing capabilities?

APPENDIX

Table 1a: Fiscal Year 2003 Budget Requests for High End Computing by Agencies Participating in the National Information Technology Research and Development program (dollars in millions)

Agency	High End Computing: Infrastructure and Applications	High End Computing: Research and Development	Total for High End Computing
DARPA	16.8	81.9	98.7
DOE/NNSA	41.4	39.5	80.9
DOE Office of Science	**98.5**	**39.3**	**137.8**
EPA	1.8	0.0	1.8
NASA	68.4	26.0	94.4
NIH	88.2	8.9	97.1
NIST	3.5	0.0	3.5
NOAA	13.3	1.8	15.1
NSA	--	31.9	31.9
NSF	**215.2**	**68.3**	**283.5**
ODDR&E	--	1.8	1.8
Total:	**547.1**	**299.4**	**846.5**

Source: NITRD National Coordination Office Fiscal Year 2003 Blue Book. The Blue Book is released in August of each year, and thus the data on FY 2003 spending and FY 2004 budget requests levels has not yet been provided to the National Coordination Office.

Note: In addition to the research and development-type activities that are counted for the data included in this table and Table 1b, many agencies devote significant funding to the purchase and operation of high-performance computers that perform these agencies' mission-critical applications.

Acronyms: DARPA—Defense Advanced Research Projects Agency, DOE/NNSA—Department of Energy's National Nuclear Security Administration, EPA—Environmental Protection Agency, NASA—National Aeronautics and Space Administration, NIH—National Institutes of Health, NIST—National Institute of Standards and Technology, NOAA—National Oceanic and Atmospheric Administration, NSA—National Security Agency, NSF—National Science Foundation, ODDR&E—Office of the Director of Defense Research and Engineering, VA—Department of Veterans Affairs.

Table 1b: Funding History from fiscal year 1992 to fiscal year 2003 of high-performance computing research and development programs at various agencies.

	FY 1992	FY 1993	FY 1994	FY 1995	FY 1996	FY 1997	FY 1998	FY 1999	FY 2000	FY 2001	FY 2002	FY 2003 (Requests)
DARPA	141.80	169.20	136.20	142.70	77.96	72.70	84.80	48.00	36.50	96.20	81.30	98.70
DOE/NNSA									113.90	168.30	75.60	80.90
DOE/SC	**73.00**	**76.20**	**84.60**	**73.10**	**84.49**	**86.00**	**90.53**	**91.90**	**84.10**	**130.30**	**126.70**	**137.80**
EPA	4.50	6.10	5.90	10.50	8.70	5.60	5.38	4.20	3.90	3.50	1.80	1.80
NASA	64.00	70.20	84.60	87.40	75.55	88.00	90.10	71.40	124.80	86.80	62.10	94.40
NIH	8.90	34.40	29.50	29.90	22.40	23.40	23.74	27.10	34.10	59.50	87.20	97.10
NIST	0.90	0.90	0.90	3.60	5.59	4.00	3.99	3.50	3.50	3.50	3.50	3.50
NOAA	1.80	9.40	9.80	2.80	3.30	4.30	4.30	8.80	13.20	12.00	15.60	15.10

NSA		40.20	32.70	28.20	29.48	30.40	26.42	24.00	31.70	32.90	41.60	31.90
NSF	**127.00**	**133.90**	**139.10**	**150.00**	**140.32**	**129.20**	**132.90**	**224.70**	**289.80**	**311.70**	**291.50**	**283.50**
ODDR&E									2.00	2.00	2.00	1.80
VA					3.00	1.00						
Totals	**421.90**	**540.50**	**523.30**	**528.20**	**450.79**	**444.60**	**462.16**	**503.60**	**737.50**	**906.70**	**788.90**	**846.50**

Source: NITRD National Coordination Office Blue Books, Fiscal Years 1992 to 2003.

Acronyms: DARPA—Defense Advanced Research Projects Agency, DOE/NNSA—Department of Energy's National Nuclear Security Administration, DOE/SC—Department of Energy's Office of Science, EPA—Environmental Protection Agency, NASA—National Aeronautics and Space Administration, NIH—National Institutes of Health, NIST—National Institute of Standards and Technology, NOAA—National Oceanic and Atmospheric Administration, NSA—National Security Agency, NSF—National Science Foundation, ODDR&E—Office of the Director of Defense Research and Engineering, VA—Department of Veterans Affairs.

Program History: Figures from FY 1992-1995 reflect the funding for the High Performance Computing Systems and the Advanced Software Technology and Algorithms Programs. Figures from FY 1996-1999 reflect the funding for the High End Computing and Computation Program. Figures from FY 2000-2003 reflect the funding for the High End Computing Infrastructure and Applications and Research and Development Programs.

Table 2: The top twenty machines of the TOP500 List of the World's Fastest Supercomputers (full list available on line at http://www.top500.org/).

Rank	Manufacturer Computer/Number of Processors	Installation Site Country/Year
1	**NEC** Earth-Simulator/5120	Earth Simulator Center Japan/2002
2	**Hewlett-Packard** ASCI Q - AlphaServer SC ES45/1.25 GHz/8192	Los Alamos National Laboratory Los Alamos, NM, USA/2002
3	**Linux Networx** MCR Linux Cluster Xeon 2.4 GHz-Quadrics/2304	Lawrence Livermore National Laboratory Livermore, CA, USA/2002
4	**IBM** ASCI White, SP Power3 375 MHz/8192	Lawrence Livermore National Laboratory Livermore, CA, USA/2000
5	**IBM** SP Power3 375 MHz 16 way/6656	NERSC/Lawrence Berkeley National Laboratory Berkeley, CA, USA/2002
6	**IBM** xSeries Cluster Xeon 2.4 GHz - Quadrics/1920	Lawrence Livermore National Laboratory Livermore, CA, USA/2003
7	**Fujitsu** PRIMEPOWER HPC2500 (1.3 GHz)/2304	National Aerospace Laboratory of Japan Japan/2002
8	**Hewlett-Packard** rx2600 Itanium2 1 GHz Cluster - Quadrics/1540	Pacific Northwest National Laboratory Richland, WA, USA/2003
9	**Hewlett-Packard** AlphaServer SC ES45/1 GHz/3016	Pittsburgh Supercomputing Center Pittsburgh, PA, USA/2001
10	**Hewlett-Packard** AlphaServer SC ES45/1 GHz/2560	Commissariat a l'Energie Atomique (CEA) France/2001
11	**HPTi** Aspen Systems, Dual Xeon 2.2GHz-Myrinet2000/1536	Forecast Systems Laboratory - NOAA Boulder, CO, USA/2002
12	**IBM** pSeries 690 Turbo 1.3GHz/1280	HPCx (UK Academic Research Center) UK/2002
13	**IBM** pSeries 690 Turbo 1.3GHz/1216	NCAR (National Center for Atmospheric Research) Boulder, CO, USA/2002
14	**IBM** pSeries 690 Turbo 1.3GHz/1184	Naval Oceanographic Office (NAVOCEANO) Stennis, MS, USA/2002
15	**IBM** pSeries 690 Turbo 1.3GHz/960	European Ctr. for Medium-Range Weather Forecasts UK/2002
16	**IBM** pSeries 690 Turbo 1.3GHz/960	European Ctr. for Medium-Range Weather Forecasts UK/2002
17	**Intel** ASCI Red/9632	Sandia National Laboratories Albuquerque, NM, USA/1999
18	**IBM** pSeries 690 Turbo 1.3GHz/864	Oak Ridge National Laboratory Oak Ridge, TN, USA/2002
19	**Atipa Technology** P4 Xeon 1.8 GHz - Myrinet/1024	Louisiana State University Baton Rouge, LA, USA/2002
20	**Hewlett-Packard** AlphaServer SC ES45/1 GHz/1392	NASA/Goddard Space Flight Center Greenbelt, MD, USA/2002